

Intellectual Property Protection Systems and Digital Watermarking

Jack Lacy, Schuyler R. Quackenbush, Amy Reibman, James H. Snyder

AT&T Labs – Research
Florham Park, NJ; Red Bank, NJ
{lacy, srq, amy, jhs}@research.att.com

Abstract. Adequate protection of digital copies of multimedia content – both audio and video – is a prerequisite to the distribution of this content over networks. Until recently digital audio and video content has been protected by its size: it is difficult to distribute and store without compression. Modern compression algorithms allow substantial bitrate reduction while maintaining high-fidelity reproduction. If distribution of these algorithms is controlled, cleartext uncompressed content is still protected by its size. However, once the compression algorithms are generally available cleartext content becomes extremely vulnerable to piracy. In this paper we explore the implications of this vulnerability and discuss the use of compression and watermarking in the control of piracy.

1 Introduction

Protection of digital copies of multimedia content – both audio and video – is a prerequisite to the distribution of this content over networks. Until recently digital audio and video content has been protected by its size. For example, audio on compact discs is encoded using PCM at 1.4 megabits per second – about half a gigabyte for a 45 minute CD. Such large quantities of data are difficult to distribute and store. Modern compression algorithms provide high-fidelity reconstruction while allowing substantial size reductions. If distribution of these algorithms is controlled, cleartext, uncompressed content is still protected by its size. However, once the compression algorithms are generally available cleartext content becomes extremely vulnerable, as is evidenced by the proliferation of illegally distributed MP3 compressed music. In this paper we explore the implications of this vulnerability and how watermarking techniques can contribute to a system strategy that protects intellectual property.

2 A Systemic View of IP Protection

The design of secure systems should be based upon an analysis of the application risks and threats. As Fig. 1 illustrates, such analysis will identify some of the risks of a particular domain. The technological net should handle many identified risks. The

Form SF298 Citation Data

Report Date <i>("DD MON YYYY")</i> 01041998	Report Type N/A	Dates Covered (from... to) <i>("DD MON YYYY")</i>
Title and Subtitle Intellectual Property Protection Systems and Digital Watermarking		Contract or Grant Number
		Program Element Number
Authors		Project Number
		Task Number
		Work Unit Number
Performing Organization Name(s) and Address(es) IATAC Information Assurance Technology Analysis Center 3190 Fairview Park Drive Falls Church VA 22042		Performing Organization Number(s)
Sponsoring/Monitoring Agency Name(s) and Address(es)		Monitoring Agency Acronym
		Monitoring Agency Report Number(s)
Distribution/Availability Statement Approved for public release, distribution unlimited		
Supplementary Notes		
Abstract		
Subject Terms		
Document Classification unclassified		Classification of SF298 unclassified
Classification of Abstract unclassified		Limitation of Abstract unlimited
Number of Pages 12		

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 074-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE 4/1/98		3. REPORT TYPE AND DATES COVERED Report
4. TITLE AND SUBTITLE Intellectual Property Protection Systems and Digital Watermarking			5. FUNDING NUMBERS	
6. AUTHOR(S) Not provided				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Information Assurance Technology Analysis Center (IATAC) 3190 Fairview Park Drive Falls Church, VA 22042			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Defense Technical Information Center DTIC-AI 8725 John J. Kingman Road, Suite 944			10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION / AVAILABILITY STATEMENT			12b. DISTRIBUTION CODE A	
13. ABSTRACT (Maximum 200 Words) Adequate protection of digital copies of multimedia content both audio and video is a prerequisite to the distribution of this content over net-works. Until recently digital audio and video content has been protected by its size: it is difficult to distribute and store without compression. Modern compression algorithms allow substantial bitrate reduction while maintaining high-fidelity reproduction. If distribution of these algorithms is controlled, clear text uncompressed content is still protected by its size. However, once the compression algorithms are generally available clear text content becomes extremely vulnerable to piracy. In this paper we explore the implications of this vulnerability and discuss the use of compression and watermarking in the control of piracy.				
14. SUBJECT TERMS Watermarking, Information Security, Crypt			15. NUMBER OF PAGES	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified		18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified		19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified
				20. LIMITATION OF ABSTRACT Unlimited

legal net will handle others. No matter how thorough the analysis, not all risks will be identified, and not all identified risks will be caught by the technological and legal nets. Ideally the system design includes the possibility of renewable security so that these residual risks do not undermine the foundations of the business.




Fig. 1.

The business model for the application is one of the strongest security mechanisms. If the system is easy to use, rich in features, support and information, and reasonably priced, why should consumers go to the black market? Designing the system with this in mind will minimize the attacks from legitimate users, most of whom are willing to play by the rules. System security should not interfere with legitimate use. We also want to design the system so that even if an attacker does break the system he cannot then use the same system to distribute that IP for his gain. Such a system should consist of:

1. a compression engine for managing music or video. This mechanism should discourage multiple compression/decompression cycles;
2. a mechanism for protecting the integrity of the content and for enforcing rights-to-use rules;
3. a flexible mechanism for licensing content and for granting various rights to consumers with appropriate credentials;
4. a secure client for accessing, rendering, playing or viewing content in a manner consistent with system policy and with the credentials or licenses associated with that content;
5. a mechanism for labeling the content to be distributed in a persistent manner. For example, the label might indicate ownership, name the distributor, identify the property or contain information about transactions involving the content.

Component 2 involves the use of cryptographic containers as in [8], [12] and [16]. The content is encrypted and perhaps digitally signed. The encryption keys are distributed via other channels using cryptographic protocols. A flexible licensing mechanism (Component 3), based for example upon PolicyMaker [3], manages these keys and governs their use [12]. Client security (Component 4) is what distinguishes the IP protection problem from the protected communications channel problem. That is,

content must be protected in the client, not just in the channel. Protection mechanisms include tamper resistant software and hardware. These techniques are discussed in [2] and [12].

2.1 Compression

As discussed earlier, compression enables the distribution of music or video over networks. For audio, the MPEG-2 Advanced Audio Coder [1] provides CD quality reproduction for most music and most listeners at a compression ratio of 11 to 1 (128 kilobits per second). Compression may also be relevant as a protection mechanism for the following two reasons.

Attackers will always have access to decompressed output. If recompression of the decompressed content results in noticeable degradation of quality, then the 2nd generation output will be of sufficiently low quality that it is not a threat to the IP owner. Of equal importance, when cleartext content is available, nothing we do to protect compressed content matters. Should controlled degradation via compression prove possible, then a solution to this cleartext audio problem would be to compress and then decompress the music as part of the mastering process. Controlled degradation via compression is an area of current research.

Because it can easily be distributed, the compressed file is the valuable commodity. It therefore makes sense to associate labels with the compressed file in a way that is persistent in the compressed domain.

3 Digital Watermarking

3.1 Overview

As stated earlier, a mechanism is needed for binding content identification to content in a persistent manner. Digital watermarking is such a mechanism. (See for example [11].) Watermarking has also been proposed as a mechanism for gating the use of content. In this case, when decisions regarding access to or use of the content are made, the mark must be retrieved in real-time and used as input in the decision-making process. No one marking algorithm is best suited for these two functions, both because of complexity issues and because different functions and different marking algorithms are resistant to different attacks. Indeed, we expect that any single album or film will be marked by a variety of different algorithms, to improve the overall resistance to attack.

3.2 System Attacks

We list several general classes of attacks against information embedded in multimedia content. The use of a watermarking algorithm for a particular application needs to be

‘sanity-checked’ against this list to determine whether or not the watermark serves any useful purpose.

If cleartext content is available and the compression algorithm is readily available to a pirate, then the pirate can generate an equivalent, unprotected, and untraceable copy of the compressed content, and bypass every protection/tracing mechanism the copyright owner might employ. Watermarking is irrelevant in this case.

One attack is forgery of identity. Whether the watermark is a point-of-sale watermark (a “fingerprint”) or pressing-plant watermark, if the input to the marking process is fraudulent, then the watermark doesn’t protect the IP.

Works distributed in versions distinguished only by different watermarks are susceptible to collusion attacks. The existence of multiple copies of a work, especially if the bitstreams differ *only* in the markings, provides a probe of the sites of the watermarks and indirectly of the marking algorithm itself. The existence of differently marked copies of a work may reduce the effectiveness of the security of the system.

The simplest collusion attack is the *bitwise XOR attack*. The attacker compares the differences between two representations of the same work, and jams differing bits to 0s or 1s. The jam pattern can be either random or – if the attacker has knowledge of the marking algorithm – one that creates a counterfeit of a legitimate mark. When a work is to be marked in multiple versions each with its own markings, the marking algorithm must be designed so that in the presence of tampering one of three conditions holds. The work should be impossible to decompress, the quality of the decompressed output should be significantly degraded, or the mark should nonetheless be recovered from the bits which were not changed by the attacker. Generally this means either that very few bits should differ from one mark to the next, or else that all of the bits in the bitstream should change when the mark changes.

Another collusion attack specific to frame-based compression algorithms can be effective against marks that extend in time through the work. In these algorithms, a bitstream is composed of frames, each representing a segment of the original signal. For a given algorithm all segments have the same duration. Given multiple versions of the signal, each marked differently, take the first frame from the first copy, the second frame from the second copy, and so on. More sophisticated versions of this attack are possible. It is difficult to see how any mark that is extended in time across several frames can survive this attack, be it a cleartext PCM watermark or a bitstream watermark. Extended watermarks should not be used for differentially marking multiple copies of a work. If an extended marking algorithm must be used for this purpose, then it should be complemented by watermarks which have a reasonable probability of recovery from bitstreams composed of fragments of watermarked streams.

System designers should think carefully before using watermarks to gate usage, since by feeding different bitstreams into the gating mechanism the attacker may be able to probe the watermark algorithm, discover mark sites and possibly generate fraudulent marks [5]. If a marking algorithm is to be used to gate usage, the algorithm should be designed in such a way that tampering with the mark should degrade the quality of the decompressed content. This suggests that the marking algorithm could beneficially be associated with the compression algorithm. We describe one such marking algorithm in section 3.5.

3.3 Desirable Characteristics of Watermark Algorithms

The following requirements are typically expected of watermarks (see also [10]):

1. *Imperceptibility*. A watermarked signal should (usually) not be distinguishable from the original signal.
2. *Information capacity*. The mark bitrate must be compatible with the rate limits imposed by the system.
3. *Robustness*. The mark must be recoverable, not only in the complete work, but also in truncated, filtered, dilated, and otherwise processed clips, in a concatenation of unrelated content, and in the presence of noise.
4. *Low complexity*. Marking schemes intended for use with real-time applications should be low complexity.
5. *Survive multiple encode-decode generations*. A watermark should survive tandem encoding-decoding.
6. *Tamper resistant or tamper evident*. It should be possible to recognize that a mark has been modified. It should not be possible to modify a mark in such a way as to create a different valid mark.
7. *Difficult to create or extract legitimate watermark without proper credentials*. In the context of the watermarking engine alone, a proper credential is knowledge of the algorithm used to insert the mark. An ideal would be a public key analogue to watermarking: hard to insert mark, easy to retrieve, hard to counterfeit.

For copyright identification every copy of the content can be marked identically, so the watermark can be inserted once prior to distribution. Ideally, detection should not require a reference because the search engine has no *a priori* way to associate the reference material with the work from which the mark is to be recovered. Not only must the watermark be short enough to be recovered in a truncated version, some means must be provided to synchronize the detection process so that the watermark can be located in the processed bitstream. Finally, any attempt to obscure the mark, including re-encoding the content, should lead to perceptible distortion.

Transaction identification requires a distinct mark for each transaction. The primary challenge of point-of-sale marking (“fingerprinting”) is to move the content through the marking engine quickly. That is, the algorithm must be low complexity. One strategy is to insert the watermark in the compressed domain, in which case mark insertion should increase the data rate very little. Watermarking algorithms designed for fingerprinting must be robust to collusion attacks.

3.4 General Mechanisms

Watermarks for compressed content fall into three categories: cleartext or original (PCM in the case of audio or video) marking, compressed bitstream marking which does not alter the bitstream semantics, and marking integrated with the compression algorithm in which the semantics of the bitstream are altered. We describe these below and discuss their advantages and limitations. We anticipate that in a well-designed system, each of these marking techniques will be used.

Cleartext PCM: We define cleartext watermarks as marks inserted in the original or during decompression into output (e.g. while writing a decompressed song to CD).

Cleartext marking embeds a data stream imperceptibly in a signal. The model for many cleartext-marking algorithms is one in which a signal is injected into a noisy communication channel, where the audio/video signal is the interfering noise [17]. Because the channel is so noisy, and the mark signal must be imperceptible, the maximum bit rates that are achieved for audio are generally less than 100bps.

Cleartext marks are intended to survive in all processed generations of the work. They are therefore well suited to identification of the work. There are two major concerns with cleartext marking. Because such algorithms (usually) compute a perceptual model, they tend to be too complex for point-of-sale applications. Second, these algorithms are susceptible to advances in the perceptual compression algorithms.

Retrieval mechanisms for cleartext watermarks fall into two classes: reference necessary and reference unnecessary. In either case the mechanism for mark recovery is generally of high complexity and is often proprietary. Further, if means for detecting these watermarks are embedded in a player, an attacker, by reverse engineering the player, may be able to identify and remove the marks. We believe that cleartext watermarks should *not* be used to gate access to content.

Bitstream Watermarking (semantic-non-altering): Bitstream marking algorithms manipulate the compressed digital bitstream without changing the semantics of the audio or video stream. Bitstream marking, being low-complexity, can be used to carry transaction information. Because the mark signal is unrelated to the media signal, the bit rate these techniques can support can be as high as the channel rate. However these marks cannot survive D/A conversion and are generally not very robust against attack; e.g. they are susceptible to collusion attacks (we describe techniques for increasing robustness to collusion in section 4.7). This type of mark can easily be extracted by clients and is thus appropriate for gating access to content; it is an example of a security measure intended primarily to “keep honest users honest”.

Bitstream Marking Integrated with Compression Algorithm (semantic altering): Integrating the marking algorithm with the compression algorithm avoids an ‘arms race’ between marking and compression algorithms, in which improvements in hiding data imperceptibly in content are undercut by and even motivate further improvements in perceptual compression algorithms. Since the perceptual model is available from the workings of the compression algorithm, the complexity associated with marking can be minimized. Integrated marking algorithms alter the semantics of the audio or video bitstream, thereby increasing resistance to collusion attacks. An example of this approach is [7], which however does not use perceptual techniques. We now present another example.

3.5 Integrating the Watermarking Algorithm with Compression

We have developed a first generation system that combines bitstream and integrated watermarking. It can be configured to support the three marking functions mentioned above. It does not include but is compatible with use of a front-end cleartext-marking

algorithm as well. We assume that the cleartext original is not available except possibly to auditors seeking to recover the watermark. In particular, the cleartext original is not available to attackers. The decompressed and marked content will generally be available to everyone.

Our method relies on the fact that quantization, which takes place in the encoder, is a lossy process. By combining mark insertion with quantization we ensure that the attacker cannot modify the mark without introducing perceptible artifacts. The fact that marking data is present is indicated by characteristics of the bitstream data. Our marking technique involves the perceptual modeling, rate control, quantization, and noiseless coding blocks of a generic perceptual coder. In MPEG AAC spectral lines are grouped into 49 “scale factor” bands (SFB), each band containing between 4 and 32 lines. Associated with each band is a single scale factor, which sets the quantizer step-size, and a single Huffman table (AAC employs 11 non-trivial Huffman tables). The coefficient for each spectral line is represented by an integer (i.e. quantized) value. In MPEG video, a block consists of 64 coefficients, and each set (termed a macroblock) of 6 blocks has an associated quantization step-size Q_p . The same Huffman table is used for the coefficients for all Q_p values. As with audio, each coefficient is represented by an integer after quantization. Because the watermarking algorithms for audio and video are similar, for consistency we use the audio terminology (scale factor) throughout when we are discussing techniques. When we discuss the results for video, we will use terminology specific to video.

Let $A = \{f_i, H_i, \{q_{ij}\}\}$ be the set of triples of scale factors f_i , Huffman tables H_i , and quantized coefficients $\{q_{ij}\}$. (Only one Huffman table is used in video.) We assume that we have selected some set of scale factor bands into which mark data will be inserted. The marking set will generally be dynamic. Let M be the set of indices associated with the set of SFB chosen for marking.

Choose a set of multipliers $\{x_i: i \in M\}$, with all x_i close to unity. Modify the triple $\{f_i, H_i, \{q_{ij}\}: i \in M\}$ as follows. Let $\{v_{ij}\}$ be the set of spectral coefficients prior to quantization, and Q_i be the quantizer for SFB i , i.e. $\forall i \{q_{ij}\} = Q_i[\{v_{ij}\}]$. Then

$$\begin{aligned} \{f_i, H_i, \{q_{ij}\}\} &\rightarrow \{f_i', H_i', \{q_{ij}'\}\}, \text{ where} \\ f_i' &= f_i/x_i \\ q_{ij}' &= Q_i'[x_i \times v_{ij}] \\ H_i' &= H_i \text{ or the next larger codebook} \\ x_i &\cong 1 \end{aligned}$$

(See [13] for changes for the slight modifications necessary for video.) Because the modification to the spectral coefficients occurs before quantization, the changes to the reconstructed coefficients will be below perceptual threshold. If this change were introduced after quantization, the change in some quantized values would be greater than the perceptual noise floor. Equivalently, an attacker who modifies the quantized values to eradicate or modify the mark will be introducing energy changes that exceed the noise floor. Because the changes in step-sizes will be small, because not all coefficients will change, and because the attacker will not have access to the uncompressed cleartext source material, the attacker will generally not be able to identify those SFB which are used for marking. Further, the increase in bit rate associated with marking should be small, and so must be monitored. A feedback mechanism similar to the one in [7] can be used to prevent modification of scale factors that would increase the bit rate significantly.

Watermark bits can be inserted in a variety of ways. Generally watermark sequences are inserted a few bits per frame. The data to be carried by the stream is typically mapped into a marking sequence prior to embedding, where the characteristics of the mapping function depend on the type of attack expected. Indeed, since there may be a wide range of attacks, the data may be redundantly mapped in different ways in the hope that at least one mapping will survive all attacks. We describe one such mapping in section 3.8.

In our system we insert the marking sequence by modifying the scale factors included at the beginning of the frame by modifying the LSBs so that they represent a sequence which contains one or more synchronization codes. Specifically, when we select a frame for watermark insertion, and a scale factor LSB does not match (0 where a 1 is indicated, or a 1 instead of a 0), we decrement that scale factor and adjust all the coefficients in the SFB accordingly. Although the watermark data can be damaged, random flipping of scale factor LSB by an attacker will introduce perceptible artifacts.

The marking sequence can be recovered by comparison to a reference or through the use of synchronization codes. Note that if synchronization codes are used, *the watermark can be recovered in the compressed domain through a lightweight recovery process*. It can therefore be used for gating access to content. Although the attacker can use the gating mechanism to probe for the watermark sites [5] and perhaps damage the synchronization codes, damage to the codes will generally produce perceptible artifacts.

3.6 Audio Results

To evaluate our audio watermarking algorithm we used AT&T's implementation of AAC. Watermark synchronization is indicated by the sequence comprising the LSB of the first 44 decoded scale factors in a long block. When the value of the LSB of a scale factor does not match the corresponding bit in the synchronization code then the scale factor is decremented and the spectral coefficients adjusted accordingly, resulting in perceptually irrelevant overcoding of the associated spectral data.

The following table shows the cost of carrying watermark data inserted into *every frame* of an AAC bitstream for a stereo signal sampled at 44.1 kHz and coded at 96 kbps. Cost is expressed as increase in bits per frame (21.3 ms of audio) and increase in rate.

An important issue for any watermarking algorithm is the quality of the reconstructed signal following an attack that obscures the watermark. We have simulated a naïve attack on this marking algorithm by zeroing all scale factor LSB, and find that this attack results in unacceptable distortion in the reconstructed audio signal.

Table 1. Increase in audio bit-rate.

	increase in bits (per marked frame)	increase in rate
Synchronization	5.2	0.25%
sync + 32 bits	9.0	0.44%

3.7 Video Results

Our baseline system for video compression uses a rudimentary perceptual model. A variance-based activity measure is used to select the quantization step-size for each macroblock as in step 3 of the MPEG-2 TM5 rate control [14]. We generate I frames every half second; all other frames are P frames. We inserted watermark data into both I and P frames, and present results taken from an average over two different 10 second sequences.

The first 44 macroblocks of a frame are used for synchronization as described in section 3.5. The next several macroblocks (100 or 600 in the Table, of 1320) of a frame carry mark bits. For each macroblock, when the LSB of the step-size Q_p does not match, Q_p is decremented. However, a dead-zone is applied to the original Q_p to ensure that zero coefficients remain zero.

We have simulated a naïve attack on this algorithm by zeroing all scale factor LSB, and find that this attack results in a perceptible 1.6dB degradation in PSNR of the reconstructed video signal.

Table 2. Increase in video bit-rate.

	increase in bits (per marked frame)	increase in rate
Synchronization	124	0.005%
sync + 100 bits	138	0.006%
sync + 600 bits	557	0.024%

3.8 Formatting Watermark Data

We said in section 3.4 that for transaction watermarking, the bits representing differently marked versions of the same content should have bitstreams which are either nearly the same or as different as possible. We have developed a simple method for formatting watermark data that is relatively resistant to XOR collusion attacks. ([4] describes an algorithm with similar intent.) Although we are using this technique for formatting watermark data for a semantic non-altering scheme, it is more generally applicable.

We assume that the set identifying data (e.g. transaction identities), one datum of which is to be formatted, can be put into a linear sequence. For example, we might uniquely mark each transaction that occurs on 31 April 1998, so we wish to identify Transaction 1, Transaction 2, and so on. Instead of representing the Nth transaction by the ordinal N, we represent it by 2^{N-1} . Assume that no further formatting of the mark data has been performed. When an attacker bitwise XORs two copies of the same content, the resulting sequence will indicate both the first transaction and the second. If the attacker sets or clears the bits identified by the XOR operation, then the resulting mark is identical to one of the original marks. If the projected bits are randomized, then the mark is invalid.

This exponential sequence is inadequate by itself as a hiding mechanism. What needs to be protected from the attacker is the location of the transitions. This can be

accomplished by permuting the bits of the sequence, possibly after XORing them with a mask. The bits of the watermark sequence can also be interleaved with other data. Finally, the watermark sequence can be redundantly inserted. These manipulations hide the transition in the watermark sequence, so that the result of an XOR of two bitstreams (which differ only in the sequences with which they are marked) appears as a random jumble of 1s and 0s.

4 Conclusion

We have discussed threat models against IPP systems, including threats posed by the existence of high-quality compression, and attacks against watermarking algorithms in particular. We have identified three classes of watermark algorithms, distinguished by the domains in which the watermark is inserted and the extent of integration with the compression algorithm. We have reviewed suggested uses for watermarking and find that a particular algorithm can be effective in some instances, ineffective in others, and compromising in yet others. There is no panacea.

We describe what we believe is the first published example of a watermarking algorithm that has been integrated with a perceptually based compression algorithm. It has the desirable property that it can be recovered in the compressed domain with a lightweight process. Although the watermark can be damaged, early work suggests that such damage will introduce perceptible artifacts.

We have also described a method for mapping watermark data into a mark sequence that is relatively robust to XOR collusion attacks.

References

1. M. Bosi, K. Brandenburg, S. Quackenbush, L. Fielder, K. Akagiri, H. Fuchs, M. Dietz, J. Herre, G. Davidson, Y. Oikawa, "ISO/IEC MPEG-2 Advanced Audio Coding", presented at the 101st Convention of the Audio Engineering Society, Nov. 1996, preprint 4382.
2. D. Aucsmith, "Tamper Resistant Software", in *Proceedings of the First International Information Hiding Workshop*, LNCS 1174, Springer-Verlag, Cambridge, U.K., May/June, 1996, pp. 317-334.
3. M. Blaze, J. Feigenbaum, J. Lacy, "Decentralized Trust Management", in *Proceedings of the 1996 IEEE Symposium on Security and Privacy*, pp. 164-173.
4. D. Boneh, J. Shaw, "Collusion-secure Fingerprinting for Digital Data", *Crypto '95*, LNCS 963, Springer-Verlag, Berlin 1995, pp. 452-465.
5. I. J. Cox and J.M.G. Linnartz, "Public Watermarks and Resistance to Tampering", *Proceedings of the Fourth International Conference on Image Processing*, Santa Barbara CA, October 1997.
6. U.S. National Bureau of Standards, "Data Encryption Standard," Federal Information Processing Standards Publication, FIPS PUB 46-1, Jan. 1988.
7. F. Hartung and B. Girod, "Digital Watermarking of MPEG-2 Coded Video in the Bitstream Domain", *Proc. IEEE ICASSP*, pp. 2621-4, April 1997.
8. "Cryptolope Container Technology", an IBM White Paper, <http://www.cryptolope.ibm.com/white.htm>.

9. Proceedings of the Fourth International Conference on Image Processing, Santa Barbara CA, October 1997.
10. International Federation of the Phonograph Industry, Request for Proposals – Embedded signaling systems issue 1.0. 54 Regent Street, London W1R 5PJ, June 1997.
11. *Proc. First International Information Hiding Workshop*, LNCS 1174, Springer-Verlag, Cambridge, U.K., May/June, 1996, pp. 207-226.
12. J. Lacy, D. P. Maher, and J. H. Snyder, “Music on the Internet and the Intellectual Property Protection Problem”, *Proc. International Symposium on Industrial Electronics*, Guimarães, Portugal, July 1997.
13. J. Lacy, S.R. Quackenbush, A.R. Reibman, D. Shur, J.H. Snyder, “On Combining Watermarking with Perceptual Coding”, submitted to *Proc. IEEE ICASSP*, 1998.
14. MPEG video committee, “Test Model 5”, ISO-IEC/JTC1/SC29/WG11 N0400, April 1993.
15. F. Petitcolas, R. Anderson, M. Kuhn, “Attacks on Copyright Marking Systems”, Second Information Hiding Workshop, 1998.
16. O. Sibert, D. Bernstein, D. Van Wie, “Securing the Content, Not the Wire, for Information Commerce”,
17. J. Smith, B. Comisky, “Modulation and Information Hiding in Images”, *Proc. First International Information Hiding Workshop*, LNCS 1174, Springer-Verlag, Cambridge, U.K., May/June, 1996, pp. 207-226.